

[Compaq]

C:\dt\negvalue

Notes on negative value of information

December 7, 1987

Marschak postulated that information could have only non-negative value; if information did not have positive value, it could be ignored, in making decisions, choosing among alternatives, or making inferences about hypotheses. Therefore it could not have negative value.

But this supposes, tacitly, a "Robinson Crusoe world," where a single actor, with a unified set of preferences, makes his choices without an audience, without judges, constituents, rivals, superiors, allies, critics: no witnesses of his behavior to make inferences as to his preferences or beliefs, to apply standards, to evaluate his performance or character, to condemn it, to predict it, or to respond to it.

Any such witnesses would have an interest in the information available to the actor as or before he made his choice--including information on the alternatives available to him, their consequences, etc.--as a basis for inferring his preferences from the choice and for making any of the other inferences or responses above. To the extent that the actor is interested in their judgments and responses (i.e., his own well-being or other goals are influenced, in part, by these) he will wish to influence their knowledge of the information available to him.

He might wish to conceal, or deceive them about, his possession of some such information, to avert inferences or responses that damage his interests. Or he might wish to make sure that they do not mistakenly attribute to him knowledge--or for that matter, ignorance--that is not in fact part of his decision-making base. In short, he has an interest in what they know, or suspect, about what he knows.

From this it follows that he has an interest in what he knows--which witnesses might find out, or guess, in various ways--which is distinguishable from the use he makes directly of this knowledge in his decisions, and which may be negative. That is, he may prefer not to know something, not to receive or correctly interpret certain information, lest this knowledge be attributed to him with damaging results. Its potential value, if any, in improving his decisions may be far outweighed by the harmful consequences to him if others knew or guessed that his decisions reflected ("revealed") this information.

This concern can lead him not only to deny or conceal his possession of information, but to take some pains, pay some price, to avoid receiving (or at least, "understanding," integrating) certain information in the first place. That information must be interpreted as having "negative value" for him: contrary to



Marschak's postulate.

The existence of negative-value information is related, in turn, to a great variety of behavior--virtually ignored in information theory and economic theory, and underanalyzed in social theorizing in general--having to do not only with secrecy and deception but with "self-deception," deliberate (and sometimes culpable) ignorance, "denial," repression and suppression, the management and sometimes falsification or destruction of memory, files and history, "anti-learning mechanisms," selective or motivated perception, avoidance of information and wilful misinterpretation or failure to "see" patterns, hypotheses, analogies. (See Orwell's description of "doublethink").

About 25 years ago, while working on my Ph. D. thesis which had to do with violations of the Ramsay-Savage-Marschak axioms of rational choice under uncertainty, I first conjectured that violations might well follow from a desire to conceal, rather than reveal, one's preferences from observers of one's choices. From this perspective, I realized that nearly all economic theorizing was, effectively, in the context of "Robinson Crusoe" decision-making, which presumed no "interested" observers. This despite the fact that "Robinson Crusoe economics" was a common pejorative phrase, supposedly referring to an earlier, outmoded form of theorizing!

In principle this should not be true of oligopoly theory, bargaining theory, game theory. Yet it was the latter that had elaborated the rational choice axioms as a basis for measuring, or really, postulating, utility payoffs. But it is a well-known aspect of real bargaining, one that any adequate theory must recognize, that participants commonly wish to conceal or lie about their real payoffs: and this could "reasonably" affect their choices, most of which are made under some observation.

Classic game theory presumes that the payoffs to all players are known correctly to all; this is not merely an unrealistic abstraction, but it specifically distracts from real activity intended to prevent this condition from applying. (This activity could lead to incorrect inferences on the utility function for a player; or it could involve violations of the axioms such that no utility function could be inferred--for perfectly "reasonable" considerations).

(I think this consideration is recognized in one case: the choice among strategies in a game, where the uncertainties of various outcomes are not represented by numerical probabilities, corresponding as they do to the "strategic" choices of a rational opponent. The minimax criterion proposed by von Neumann and Morgenstern for the case of the zero-sum two-person game leads to the violation of the von Neumann-Morgenstern axioms for "risk prospects," gambles with no "game" or adversarial aspects. But I am



suggesting that the "game" context come close to being a norm, for everyone but a literal Robinson Crusoe. At least, it is the norm for both political and organizational behavior, which encompasses most of the behavior (other than consumer choice) in which economists and political scientists are actually interested in.)

At any rate, it was in the course of these reflections that I first conceived, a quarter-century ago, of the existence and significance of negative-valued information. (This was not reflected in my thesis or any other published writing; I have collected notes on the notion ever since). Indeed, the first examples I thought of did not involve observers, or game interactions, but another dimension of experience ignored by economic theorizing: feelings, as distinct from preferences and choices.

Certain information might affect our "happiness," without changing any of our choices, even our potential choices (thus, without having any "value" in the Marshak sense (?)). We might not be able to do anything about this information, to use it in any way: yet it could cause pain, anxiety, guilt, or the reverse. Thus, the information that a dead mother had not loved us (to take an example at random); that a former spouse had been unfaithful; that a former, respected boss had participated in a criminal conspiracy; that a war in which one had patriotically participated had been aggressive: all this could be information one would much rather not have.

There might be circumstances in which one could use this knowledge, by analogy (especially the last three), but even then, not necessarily to an extent that would compensate for the loss of confidence, trust, self-esteem, and happiness associated with such discoveries. Therefore one could predict mental and other activity aimed at "warding off," averting such knowledge, guided by hints or suspicions that it might be coming toward us.

What came to me yesterday, Sunday December 6, was the thought that another whole class of negative-value information arises in a social situation somewhat distinct from the "game" context of strategic, adversarial interaction: the realm of obligations, duties, rules of conduct. (Again, these are concepts almost entirely foreign to economic theory, though familiar in political theory and ethics). Economic theory (and systems analysis, which derives from it) would try to model such constraints in terms of personal utilities, preferences: clumsily and inadequately. But by abstracting from the social context of these notions, this misses, among other things, the interest one might have in evading such obligations in some cases, or avoiding a judgment that one has failed to meet them, by manipulating the reality, or the impression that witnesses have, of the information available at the time one acted.



Obligations may have to do directly with potential information. One may be obliged to pass it on; to inform, warn, alert someone. If they suffer in their own duties or welfare from a failure to receive such information from you--who is generally placed and expected to have such information and to pass it on--one of the few excuses available is that you did not have the information yourself. Even if you "should" have had it, and your failure to have it is in some reasonable sense culpable, the fact that you did not know it is almost always a major extenuating consideration if not a total excuse (except for Navy captains whose ships run aground or have a collision). In nearly every context, contrary to a familiar but misleading precept, Ignorance is an excuse. Therefore, ignorance can be desirable, and deliberately achieved.

The same applies to obligations to act in certain ways on information: to carry out a duty; to show concern for others' feelings or welfare; to protect someone else's rights or interests.

A failure to perform as obliged, or to report as obliged, may be the result of a mistake, unmotivated and unintended, inadvertent ignorance. In this case, the interest in demonstrating that it was such is simply an effort, after the event, to reduce or avoid the condemnation or other social consequences of the failure.

But, as everyone knows, this effort may be disingenuous. It is very commonly the case that obligations do not correspond exactly to our personal preferences in a given situation. They may have been imposed, involuntarily. And even if we accepted them voluntarily--as a condition of employment or membership, or as constraints that we agree "should" be followed ("most of the time, by most people") we may not prefer to follow them in this particular case. We may, then, want the excuse of ignorance, and hence even the reality of ignorance, to permit us to escape condemnation or punishment for evading the obligation.

"What did the President know and when did he know it?" Two Presidencies have hung on that question. (It was not asked of Lyndon Johnson about Vietnam, because people thought they knew the answer--they thought he had been misled by the military--but the Pentagon Papers gave a different answer, for him and three others). Did he have "guilty knowledge" of crimes committed by his subordinates: which should have led him to report these to the Attorney General and to Congress (and the public), and to take immediate action to stop them, reverse their effects, compensate victims, and take precautionary measures against future occurrence. (Assuming his knowledge did not derive from having ordered these crimes himself, which would be directly culpable).

The presumption that "ignorance is an excuse"--virtually a complete exoneration from impeachment or conviction of a President--was demonstrated by the criticality of the search for the "smoking



gun" that would reveal unequivocal knowledge, and by the Presidential efforts to deny and conceal the existence of such evidence.

Likewise, both in the Vietnam case and in Irangate (and in Cuba I and Cuba II--but not, so far as we yet know, in Watergate) it was regarded as critically important to conceal from the Congress and public the existence of controversy within the inner circle of Presidential decision-making. And in all these cases, most dissent was successfully concealed for prolonged periods. (In the case of Cuba II, the Cuban Missile Crisis, for 25 years: including the fact that on Saturday, October 27, 1962, it was the President who was the dissenter from a consensus that ruled his basic decision that day not to accept, immediately, the trade with the Turkish missiles that Khrushchev proposed).

The purpose is to conceal the fact that the President had available opinions or hypotheses or recommendations (in the above case, his own!) that put in question the desirability of the course publicly chosen. The President wishes to give this course the authority of unanimous recommendation among all those who have available to them the information available to the President: which is, of course, "much greater and more reliable than that possessed by anyone else," and which cannot, unfortunately, be shared with others lest it become known to foreign adversaries.

The fact that some who had this same base of information, or close enough to it to be included among the President's advisors, interpreted it differently or saw it favoring a different course of action is very subversive of this authority, validating the doubts and dissent of outsiders, who would be encouraged by this information to press their own case to others and to the President. The President wants to be able to say, "If you knew what I know, and had the same broad values, the sense of national interest, you would choose as I have; everyone advising me, who has this information, agrees." When this is untrue, great effort is taken, generally with success, to conceal it, and especially to conceal Presidential awareness of it.

Eisenhower, in fact, was famous for not wanting to know such disagreements among his top subordinates; he wanted the NSC process to arrive at compromises that obscured such discordance before "decisions" arrived on his desk. (But he was perfectly capable of letting a subordinate like Dulles appear responsible for an unpopular or risky decision--e.g., refusal to finance the Aswan Dam for Nasser--even when he had personally authorized it. His decision to take responsibility for Powers' U2 flight was so exceptional that most people, including Khrushchev, did not believe him--though he was telling the truth!)

Hence the subversive potential of--and the extreme resentment generated by--the RAND letter I signed with five others which



called for unilateral withdrawal from Vietnam within one year, having identified the signers as Vietnam analysts all of whom had had access to classified information and had been trusted by the Administration to conduct classified research for it. The message was: You don't have to be ignorant of classified, official information ("like Abby Hoffman") to favor unilateral withdrawal, as the Administration had always implied and the establishment media had faithfully echoed.

Likewise, the damning effect of David Stockman's revelations that accurate predictions of the deficit-increasing effects of Reagan's tax and budget policies had been presented to the President early on, and with the authority of the same person constructing the programs. As with the revelation in the Pentagon Papers that military and intelligence predictions of the costs and forces needed in Vietnam and the stalemate to be expected with existing programs, the effect of the knowledge is to fasten personal responsibility for the programs squarely on the President--"He was not ignorant, or at least, he need not have been, given the information and opinions available to him"--at the same time as raising questions about the reasonableness and promise of his choice: "How could he?" Similar questions are raised by the belated information on the warnings that preceded the Challenger disaster: and indeed, many other catastrophes that decision-makers would prefer to be perceived by the public to have been unforeseeable, or at least, unimagined.

The answer to the question, "How could they?" is another story. The point here is that decision-makers have a great stake in precluding that question from being raised widely in society, in particular among their general supporters, and therefore are aware of the dangers of foreknowledge (as well as the possible benefits, known to information theorists and everyone else.) The resolution of this tension between the need for foresight, to reduce the likelihood of catastrophe, and the danger of public knowledge of one's foresight in case catastrophe is not avoided, relies on secrecy and lies, the manipulation of public knowledge about the decision-making process that preceded the disaster.

Thus, one of the Watergate tapes--taped by Ehrlichman himself in his office--shows Ehrlichman removing from a collection of files relating to the "plumbers" the one in which he authorized Young and Krogh to burglarize Dr. Fielding's office, with his handwritten notation on the margin, "If done with your assurance that it will not be traceable." He took this away from Young, to destroy it himself, with the words, "This shows a little too much foreknowledge." (Unfortunately for Ehrlichman, Young had taken the precaution of xeroxing this authorization for his own collection, with which he bought freedom from prosecution).

Likewise, Admiral Poindexter claimed--almost surely falsely--to have taken it on himself to keep the President ignorant of the



diversion of funds from Iran to the contras, successfully shielding him--not legally, but politically--from impeachment proceedings. He also took responsibility for destroying a Presidential finding that explicitly authorized trading arms for hostages: something the President still says (interview on Dec. 5 with network anchors) he did not mean to do and did not do (though some advisors were right in warning that he would be perceived by others as having done this). Supposedly Poindexter destroyed this without asking or telling the President, because, in effect, "it showed too much foreknowledge" ("It didn't reflect the President's real thinking, and could be misunderstood" even though the President had signed it. In other words, if one assumed that the President read what he signed (not, in fact, a necessary assumption, especially for this President) this document would have undermined the President's claim that he had never perceived his transactions in these terms, even though such a perception would seem virtually inescapable.

The President's forgetfulness and fuzzymindedness and apparently extraordinary powers of wishfulness and denial in managing his own awareness (his willingness to reflect wishfulness and denial in his public expressions is beyond any doubt) all make it impossible to know or prove beyond a reasonable doubt that his decision-making reflected or revealed any particular set of premises: goals, values, expectations, beliefs. Not only does this protect him from conviction of culpability for a particular decision, but it makes confident prediction impossible.

Does the apparent change in his position on accords with the Soviets reflect unforeseen flexibility, or, as he would now like to say, was the possibility implicit in his position all along? There is no way to be sure; or to guess how much further this could go? (For the record, I did come to believe that he did have core commitments--e.g., to the continuation of an arms race, modernization, and to SDI--and this has not yet been disproven, though my confidence is shaken. Could the departure of Weinberger, Perle and Gaffney, Regan and Casey, leaving the field to Nancy, have made a big difference?)

At issue throughout these matters was the President's obligations to "take care that the laws were faithfully upheld"--thus, to avoid giving illegal orders, and if illegal actions were taken, to assure that the Attorney General investigated and prosecuted them; to inform Congress of covert actions proposed or undertaken; to certify to Congress (via a "finding") that national security interests required a particular covert action.

Let B be the person(s) that A seeks to reassure by accepting an obligation or promise, as in a contract or unilateral declaration. If A's preferences over possible actions corresponded exactly to B's preferences for A's actions--if A wanted to do in all circumstances exactly when B wanted A to do--then undertaking an obligation to this effect would not influence A's behavior, nor



would it be necessary except to persuade B that this state of affairs existed. Nor would it give rise to any inclination on A's part to avoid or deny or repress particular information. The latter would also be true if undertaking the obligation affected A's payoffs and preferences (e.g., by penalizing him in various ways if he broke the promise) so as to make it true that he now preferred to do what he was obliged to do, what B wanted him to do.

But when, despite the existence of an obligation (initiated by A or imposed by B or by others), there exists a divergence between what A would personally, subjectively, prefer to do and what he is obligated to do--what he "should" do in light of his obligation, or what B would prefer that he do, and expects and demands that he do in light of the obligation, under certain circumstances, which include his subjective knowledge and understanding of various circumstances--then A has an interest in his own ignorance or misunderstanding of certain circumstances, in ways that permit him--without exposing himself to sanctions for violating his obligations--to follow his own preferences. I.e. he is motivated to misunderstand or misperceive the actual circumstances (or the actually available information) in ways that negate any import of his obligations that he "should" do something other than what he is doing, which is what he personally wants to do in the actual circumstances.

[did you follow that?] [these are just notes]

The obligation in question need not be made to other persons, and it need not affect the welfare of other persons (though both of these conditions are typical). It may be made to "oneself." But in this case, the recipient and/or monitor of the obligation and its implementation is best conceptualized as a separate part of the self or personality from that part that calculates "personal, perhaps selfish but at any rate ego-determined preferences": e.g., the ego-ideal or superego, vs. the ego.

The concerns that actor is "supposed" or obligated to take into account appropriately may involve the welfare of other people, or they may involve deontological rules, or his own (and others') welfare in the long-run, or in spheres other than the immediate one. What can give rise to double-think, to deliberate ignorance or self-deception, can be a divergence between the demands of one's own conscience, and what one would prefer to do if the voice of conscience, or of the long-run, or of "constituents" or victims, were silent in this case: or could be silenced, or modified, or muffled.

The point is that the latter manipulations are not outside one's set of capabilities. And this is still more true when "one" has a bureaucratic structure, which can be designed and operated precisely to make a particular actor ignorant of, or misinformed about, certain data available to other parts of the bureaucracy.



(Indeed, when an individual behaves likewise, it may well be convenient to conceptualize his inner information-handling process as if it were bureaucratic.)

Some information may not change our preferences within a given set of actions, yet so lower our expectations for each of these actions, or specifically for the best of them, as: a) to make us much less happy, confident, hopeful (the earlier case of "negative-value information"); and (b) to make this set of actions look highly inadequate, challenging us either to include certain actions that had earlier been excluded because of certain drawbacks (e.g., "revolution," foreign intervention, costly or painful adaptations: as in the case of Haitian elections or the redesign of electric power grids or the recall of products: see below), or to search in what may be a costly, painful way, for new alternatives or hypotheses and understandings (where the acknowledged need to do this may call into question our past and current qualifications for leadership or office, making us vulnerable to rivals).

The very stimulus to do this reexamination or search may be seen as the positive value of the otherwise-unhappy information. "After all," if the new decision process does not seem, on balance, to promise net benefit, it need not be undertaken; the unwelcome information can be "ignored," so its value cannot be negative. (Marschak's point). But the result of ignoring it, once it has become public (known to others) (or to one's own "conscience") may be reproach, removal from decision-making power, loss of self-respect, or other sanctions that make this awareness "no-win," negative on balance, worsening one's prospects whether one responds to it or does not. Again, this flows from the non-Robinson Crusoe context of (nearly all) actual decisions, and the differences in the preferences of different people as to what one will do and between these and various obligations and standards governing what one "should" do.

The actual stimulus for these notes was the newspaper stories on December 6 on the Haitian elections and the newly-recognised (and still under-investigated) risk that low electric fields, as in common household appliances, might cause cancer. In the former case especially, strong warning "signals" were ignored by most parties in Haiti, who wanted to believe that the elections--their "only hope" --would take place and would register the public will (the only alternatives being extremely dangerous and unwelcome). Thus the actual violence and intimidation came as a "surprise." Similarly, the challenge (which, contrary to Haitian case, could more easily be dismissed as "incredible") of the growing indications that the entire electrical system of the world might have to be reshaped to reduce its contribution to cancer.

The Reagan Administration, and others', reaction to AIDS, in terms of societal investigation (or individual test!) would be another example. (Or the problem of herpes; or other STDs; or, for



sexual fundamentalists, the problem of teenage pregnancies, or the entire world population question (see US refusal to fund population control programs that tolerate or encourage abortion).

Thus, knowledge that one has AIDs virus can only permit you to observe your own, or societal, obligations to other people, to protect them from your own pleasure-seeking activities, at vast cost to your own pleasure. So far, ignorance of this fact (failure to take the test) is not regarded as highly culpable, nor is taking some risk (for a member of a "low risk group") while ignorant. The standard of what one "should" do is already clear. But even if taking the test were costless in terms of effort and money, one should expect great resistance to this act of "finding out." As of now, the psychological penalty from failing to meet this social obligation does not overweigh the possible cost of observing it.

The same phenomena could arise from conflicts of obligations, as well as from the conflict between obligations and personal preferences. Thus, they could arise in the Milgram experiments, where (apparently weaker) obligations to protect the welfare of strangers conflict with an obligation to fulfill one's promise or contract to carry out the conditions of the experiment. (In reflecting on these results, I realized that every obligation is in some degree "an obligation to oneself," to demonstrate that one is a responsible, reliable person who intends to keep contracts, promises, obligations, and has the self-control to do this. The Milgram experiment shows how this can generate conflicts with obligations to protect others. This in turn can lead to denial, wishfulness, self-deceptions and the other phenomena considered here.